

REMARKS

This Amendment is filed in response to the Office Action mailed on August 20, 2008. All objections and rejections are respectfully traversed.

Claims 6-9, 11-17, 19-20, and 23-43 are currently pending.

Summary of Examiner Interview

Applicant respectfully thanks the Examiner for the Interview held on January 21, 2009, where the claim amendments to claim 6 were discussed. The Examiner stated the revised claim is allowable over the cited art, but a new search is necessary.

Request for Interview

The Applicant respectfully requests a telephonic interview with the Examiner after the Examiner has had an opportunity to consider this Amendment, but before the issuance of the next Office Action. The Applicant may be reached at 617-951-3067.

Claim Rejections – 35 USC §101

At paragraph 3 of the Office Action, claims 27, 35, and 50 were rejected under 35 U.S.C. § 101, as being directed to non-statutory subject matter.

Applicant has amended the claim to state a computer readable storage medium containing program instructions executed by a processor. The processor is hardware and

the computer readable storage medium can be a flash drive or a cd-rom for example.

Therefore, the claims are believed allowable over the §101 rejection.

Claim Rejections – 35 USC §112

At paragraph 4 of the Office Action, claim 17 was rejected under 35 U.S.C. §112, second paragraph, for insufficient antecedent basis.

Applicant has amended the claim to overcome the rejection, and believes the claim is allowable over §112 rejection.

Changes to the Specification

The new paragraphs added to the specification from US Application Serial No. 10/027,020 (Attorney Docket No. 112056-0006). The new paragraphs and the new Figure 7 do not add new matter because US Application Serial No. 10/027,020 was incorporated by reference and was filed on the same day as the pending application.

Claim Rejections – 35 USC §103

At paragraph 6 of the Office Action, claims 6-8, 12-17, 19, 27-29, 42-43, and 49- were rejected under 35 U.S.C. §103 as being unpatentable over Nishanov et al., US Patent Application Publication No. 2003/006782, hereinafter Nishanov in view of Wang et

al., US Patent No. 6,470,382, hereinafter Wang, and in further view Kawashima, US Patent No. 5,894,588, hereinafter Kawashima..

The present invention, as set forth in representative claim 6, comprises in part:

6. A method of claiming ownership of a plurality of disks by a network device of a plurality of network devices in a network storage system, comprising:

writing ownership information to a predetermined area of each disk, wherein the predetermined area of the disk is sector 0 on the disk and the ownership information stored in sector 0 is definitive ownership data for determining ownership of the disk;

setting a small computer system interface (SCSI) reservation tag for each disk to a state of network device ownership to provide a two part indicia of ownership for each disk, where the two part indicia of ownership are both written to each disk;

creating a table on each network device in the network storage system;

identifying all disks owned by each network device using ownership information written to the predetermined area of each disk of the plurality disks and, for each identified disk, if a mismatch occurs between the ownership information on the predetermined area of the disk and the ownership defined by the SCSI reservation tag, then *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag*;

in response to identifying, storing entries in the table, wherein each entry identifies an owned disk of the network device storing the table;

identifying, by a second network device, all disks owned by a first network device in response to a failure of the first network device, wherein each network device can read ownership information of each disk;

setting a SCSI release tag for each disk owned by the first network device, in response to the failure of the first network device, to transfer the disk to an unowned state; and

removing ownership information stored in the predetermined area of each disk owned by the first network device to complete transferring each disk into the unowned state.

By way of background, Nishanov discloses a system for protecting a node's exclusive access to a storage device. Each node is assigned an initiator ID and the initiator ID is mapped to a reservation key in a key table, wherein the reservation key table is stored on each disk. The reservation key includes three fields, which are a version, ReserveID field, and an OwnerID field. Each reservation key is mapped to a reservation type in a reservation table, wherein the reservation table is stored on each disk. The reserve type is configured with a persistent reservation. Each storage device is always registered. When a challenge occurs for ownership of the storage device, a new node writes ownership information to the ReserveID field. The owner then has time to win the challenge by writing its ownership information in the ReserveID field. If the owner writes its ownership information to the ReserveID field before the time is up, then the new node (challenger) loses the challenge. If the owner does not write its ownership information to the ReserveID field before the time is up, then the new node wins the challenge and the new node writes its ownership information into the OwnerID field.

Wang discloses a table on each server stating disk ownership information.

Kawashima discloses storing ownership information in sector 0.

Applicant respectfully urges that Nishanov, Wang, and Kawashima, taken alone or in any combination, does not teach or suggest Applicant's claim novel *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag, ... identifying, by a second network device, all disks owned by a first network device in response to a failure of the first network device, wherein each network device*

can read ownership information of each disk, setting a SCSI release tag for each disk owned by the first network device, in response to the failure of the first network device, to transfer the disk to an unowned state; and removing ownership information stored in the predetermined area of each disk owned by the first network device to complete transferring each disk into the unowned state. In further detail, Applicant's claimed invention uses two part ownership identification method. The first part of this ownership method is writing ownership information to a predetermined area of each disk. Within the system, ownership information written to the predetermined area of each disk is the definitive ownership attribute. This predetermined area of the disk can be any known and constant location on each of the disks and is usually sector 0 on each disk. The second part of the ownership method is setting of a SCSI reservation to allow only the disk owner to write to the disk. This use of a SCSI reservation allows other servers to read the ownership information from the disks. Furthermore, the SCSI reservation does not require storing read-only access permission for each network device because each network device not owning the disk can already read the disk. Additionally, the ability to change the SCSI reservation tag to match the ownership information stored in the predetermined area of disk allows a storage server to configure the disks into the appropriate RAID groups and or volumes. The ownership information stored on the storage device (sector 0) is the controlling ownership information on the disk without requiring the owned network device to send a second SCSI reservation tag. When a mismatch occurs between the ownership information stored in the predetermined sector and the SCSI reservation, the SCSI reservation is changed to show the same owner that is in the predetermined sector of the disk. Furthermore, each network device stores a table listing the different disks that a particular network device owns.

When a first network device fails, a second network device changes the ownership to an unowned state in both sector 0 and in the SCSI reservation tag. The second network device then writes its ownership into sector 0 and the SCSI reservation tag to claim ownership of the disk.

There is no disclosure in Nishanov of changing to unowned state to release a disks ownership. Nishanov changes ownership by using a second SCSI reservation to challenge ownership. Furthermore, neither Wang nor Kawashima discloses changing ownership to an unowned state.

Accordingly, Applicant respectfully urges that the combination of Nishanov, Wang, and Kawashima is legally insufficient to make obvious the present claims under 35 U.S.C. §103 because of the absence of the Applicant's claimed novel *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag, ... identifying, by a second network device, all disks owned by a first network device in response to a failure of the first network device, wherein each network device can read ownership information of each disk, setting a SCSI release tag for each disk owned by the first network device, in response to the failure of the first network device, to transfer the disk to an unowned state; and removing ownership information stored in the predetermined area of each disk owned by the first network device to complete transferring each disk into the unowned state.*

At paragraph 7 of the Office Action, claims 20, and 23-25 were rejected under 35 U.S.C. §103 as being unpatentable over Nishanov, Wang, and Kawashima, in view of Carlson et al., US Patent Application Publication No. 2003/0093501, hereinafter Carlson.

The present invention, as set forth in representative claim 20, comprises in part:

20. A network storage system comprising:

one or more switches interconnected to form a switching fabric;

a plurality of disks, each of the disks connected to at least one of the switches, each disk storing a first ownership attribute to a predetermined area of a disk and each disk associated with a second ownership attribute in the form of a small computer system interface reservation, wherein the predetermined area of the disk stores definitive ownership data for determining ownership of the disk and the small computer system interface reservation allows other network devices to read the ownership attribute from the disks; and

one or more network devices, interconnected with the switching fabric, each of the network devices being configured to own a predetermined set of disks of the plurality of disks through use of the first and second ownership attributes, wherein each network device identifies all disks owned by the network device using ownership information written to the predetermined area of each disk of the plurality disks and, for each identified disk, if a mismatch occurs between the ownership information on the predetermined area of the disk and the ownership defined by the SCSI reservation tag, then *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag* and each network device is configured with a table and to store entries in a table, wherein each entry identifies an owned disk of the network device storing the table.

By way of background, Carlson discloses a Storage Area Network (SAN) where storage devices are interconnected by switches to form a fabric. *See* paragraph 0039.

Applicant respectfully urges that Nishanov and Carlson taken alone or in combination do not teach or suggest Applicant's claimed novel *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag*. As stated above, Applicant's invention, stores ownership information in a predetermined sector and

persistent reservation on disk. Also, each network device stores entries in a table on the network device, where each entry identifies each disk that the network device owns.

Nishanov requires sending a second persistent reservation to maintain ownership when the two reservations are not the same. There is no disclosure in Nishanov of *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag*. Applicant's invention uses the predetermined area (sector 0) as the controlling ownership information.

There is no disclosure in Carlson, Wang, and Kawashima *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag*.

Accordingly, Applicant respectfully urges that the combination of Nishanov, Wang, Kawashima and Carlson is legally insufficient to make obvious the present claims under 35 U.S.C. § 103 because of the absence of the Applicant's claimed novel *using the ownership information written to the predetermined area of the disk as definite ownership data without requiring the owned network device to send a second SCSI reservation tag*.

At paragraph 8, of the Office Action, claim 26 was rejected under 35 U.S.C. §103 as being unpatentable over Nishanov, in view of Carlson, and in further view of Jaskiewicz et al., US Patent Application Publication 2003/0061491, hereinafter Jaskiewicz.

Applicant respectfully notes that claim 26 is a dependent claim that depends from an independent claim which is believed to be in condition for allowance. Accordingly, claim 26 is believed to be in condition for allowance.

At paragraphs 9 of the Office Action, claims 28-42 and 45-48 were rejected under 35 U.S.C. §103 as being unpatentable over Brunelle et al., US Patent No. 6,654,902, issued on Nov. 25, 2003, hereinafter Brunelle, in view of Nishanov, Wang, and Kawashima.

The present invention, as set forth in representative claim 28, comprises in part:

28. A method for a network device to manage ownership of one or more storage devices in a network storage system, comprising:

reading ownership information from a predetermined area of each storage device, wherein the predetermined area of each storage device is sector 0 on the disk and the ownership information stored in sector 0 is definitive ownership data for determining ownership of the storage device;

in response to reading the ownership information, creating an ownership table that stores entries where each entry identifies a storage device owned by the network device, wherein the ownership is stored within the network device;

reading a small computer system interface (SCSI) reservation tag from each storage device, wherein the SCSI reservation tag allows other network devices to read the ownership information from each storage device;

comparing the SCSI reservation tag to the ownership information of the same storage device and, if there is not a match, changing the SCSI reservation tag to match the ownership information;

configuring the one or more storage devices identified in the ownership table into at least one volume for use by the network device;

identifying, by a second network device, all storage devices owned by a first network device in response to a failure of the first network device, wherein each network device can read ownership information of each storage device;

*setting a SCSI release tag for each storage device owned by the first network device, in response to the failure of the first network device, to transfer the storage device to an unowned state; and
removing ownership information stored in the predetermined area of each storage device owned by the first network device to complete transferring each storage device into the unowned state.*

By way of background, Brunelle discloses a way of using standard small computer system interface (SCSI) persistent reservations with I/O barriers. The American National Standards Institute (ANSI) has standardized a number of SCSI Persistent Reservation commands, such as *Persistent Reserve Out*. See col. 1, lines 28-41. Brunelle describes issuing two of these commands to assign ownership to storage devices. See col. 5, lines 60-67. The first *Persistent Reserve Out* command includes a key describing a particular node owning the device. See col. 6, lines 38-48. The second *Persistent Reserve Out* command includes a parameter specifying an access type, such as “write exclusive read only.” See col. 5, lines 65-67 and col. 6, lines 48-54. Additionally, Brunelle describes writing registration information each time a node is initialized or changed.

Applicant respectfully urges that Brunelle Nishanov, Wang, and Kawashima, taken alone or in combination, do not teach or suggest Applicant’s claimed novel *identifying, by a second network device, all storage devices owned by a first network device in response to a failure of the first network device, wherein each network device can read ownership information of each storage device, setting a SCSI release tag for each storage device owned by the first network device, in response to the failure of the first network device, to transfer the storage device to an unowned state; and removing ownership information stored in the predetermined area of each storage device owned by*

the first network device to complete transferring each storage device into the unowned state. As stated above, in Applicant's invention, when a first network device fails, a second network device changes the ownership to an unowned state in both sector 0 and in the SCSI reservation tag. The second network device then writes its ownership into sector 0 and the SCSI reservation tag to claim ownership of the disk.

As stated above Nishanov does not *identifying, by a second network device, all storage devices owned by a first network device in response to a failure of the first network device, wherein each network device can read ownership information of each storage device, setting a SCSI release tag for each storage device owned by the first network device, in response to the failure of the first network device, to transfer the storage device to an unowned state; and removing ownership information stored in the predetermined area of each storage device owned by the first network device to complete transferring each storage device into the unowned state.* Additioanlly, Brunelle does not the stated elements above.

Accordingly, Applicant respectfully urges that the combination of Nishanov, Wang, Kawashima, and Brunelle is legally insufficient to make obvious the present claims under 35 U.S.C. §103 because of the absence of the Applicant's claimed novel *reading ownership information from a predetermined area of each storage device, wherein the predetermined area of each storage device is sector 0 on the disk and the ownership information stored in sector 0 is definitive ownership data for determining ownership of the storage device, in response to reading the ownership information, creating an ownership table that stores entries where each entry identifies a storage*

device owned by the network device, wherein the ownership is stored within the network device.

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims.

The Applicant respectfully solicits favorable action.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

/Shannen C. Delaney/
Shannen C. Delaney
Reg. No. 51,605
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500